When most people hear “Micro SaaS,” they picture code. User authentication, subscription billing, a dashboard with a logo they spent too long designing, an API that took three weeks to figure out. The technical barrier feels significant — significant enough that most non-developers never start.

Here’s the reframe that changes everything: most successful Micro SaaS products are not primarily software. They are primarily information. The software is a wrapper. The database is the business.

Once you understand that, the path to building one looks completely different — and significantly more accessible to anyone willing to think carefully and work consistently over time.

Look Beneath Most SaaS Products

Take almost any successful niche SaaS product and look at what it actually is underneath the interface.

Job boards: a database of open positions, filtered and searchable. Lead generation tools: a database of companies, enriched with contact and behavioral data. Real estate platforms: a database of properties, transactions, and market data. AI tool directories: a database of products, categorized by function and use case. Market research tools: a database of industry data, organized for fast retrieval.

In each case, the asset is not the interface. The interface is just how users access the asset. The asset is the organized, maintained, continuously updated information underneath.

This matters because building and maintaining a database is a fundamentally different kind of work than building software. It requires research, judgment, categorization, and patience — not necessarily a computer science degree or a development team. It requires the kind of systematic, curatorial thinking that one person with the right process can actually sustain.

And AI has changed what one person’s process can produce.

Prompts As Workers, Not Just Tools

Most people use AI prompts to create content. Write this, summarize that, draft this email. That’s useful, but it’s not the most interesting application for someone trying to build an intelligence asset over time.

The more powerful use is deploying prompts as workers inside a pipeline.

Imagine you’re building a database of SaaS companies in a specific niche. Rather than one prompt that tries to do everything, you build a sequence where each prompt handles one function:

Worker one researches the company — founding date, product category, target customer, pricing model. Worker two writes a plain-language summary of what the product does and who it’s for. Worker three categorizes it against your taxonomy — industry, company size served, deployment type. Worker four assigns tags based on features and integrations. Worker five identifies the two or three closest competitors and notes the key differentiation. Worker six scores the opportunity — market size, competitive density, growth indicators.

Now you have an intelligence pipeline. Each company that moves through it emerges categorized, summarized, tagged, cross-referenced, and scored. One person oversees the process. The prompts do the analytical labor.

See also The Art of Writing Human-Sounding AI Content

This is not magic. Each worker prompt needs to be built carefully and refined over time. The taxonomy needs to be designed before the categorization can be consistent. The scoring criteria need to be defined before the scoring means anything.

But that’s exactly the kind of thinking that creates a durable asset — because that thinking is what makes your database useful rather than just large, and it’s what makes it hard to copy even after you’ve made it public.

The Question Is More Important Than the Technology

Before building anything, there’s a prior question that determines whether the whole project is worth pursuing.

What valuable question does this answer?

The best intelligence systems are built around one question that a specific group of people genuinely needs answered — and currently can’t answer easily without doing a lot of work themselves.

Which AI tool should I use for this specific task? That question is worth a well-organized directory to thousands of people daily. Which local businesses in my city are underserving their customers online and likely to need marketing help? That question is worth a lead intelligence platform to an entire category of agency owners. Which newsletters in my niche are growing, what are they charging for sponsorships, and what subject lines perform best in this space? That question is worth a newsletter intelligence database to anyone trying to grow or monetize their own list. Which keywords in my industry have real search volume but weak existing content competing for them? That question is worth an SEO intelligence platform to every content marketer operating in that space.

Notice what all of these have in common. They’re not asking for content. They’re asking for organized intelligence that helps someone make a better decision faster than they could on their own. That’s the product. The database is how you deliver it. The interface is how they access it.

The question matters more than the technology because it determines whether you’re building something people actually need or building something technically interesting that nobody pays for.

Patience Is the Real Competitive Moat

Most people who attempt to build intelligence systems give up too early. Not because the idea was wrong. Because the value is invisible at the beginning, and invisible value is easy to mistake for no value.

A database with 100 records is a curiosity. A database with 500 records is a proof of concept. A database with 1,000 records starts to feel like something — but it’s still not particularly hard to replicate.

The moat appears somewhere between 10,000 and 100,000 records. At that scale, the database is genuinely difficult to reproduce. The time investment required to catch up becomes a real barrier. The cross-referencing and interconnection between records starts creating value that no individual record carries on its own. The patterns become visible that only emerge at scale.

See also How I Build AI Workflows Instead of One-Off Prompts

Nobody sees compounding at the beginning because there’s nothing to see yet. The work in month one looks identical to the work in month twelve — collecting, categorizing, enriching, updating. But the asset those two months of identical work are contributing to are completely different things.

This is where patience becomes a genuine competitive advantage rather than just a virtue. Most builders quit at 500 records because nothing has happened yet. The person who reaches 50,000 records has built something that the person who quit at 500 can never catch up to quickly — not because they were smarter or better funded, but because they kept working while everyone else stopped.

That asymmetry is rare in business. It’s one of the most interesting things about intelligence systems as a model.

Start Manual. Automate Later.

There’s a strong temptation when building anything AI-adjacent to automate everything immediately. Build the pipeline first. Set it running. Let it fill the database while you sleep.

This almost always produces a database full of low-quality, inconsistently categorized, poorly structured records that require more work to fix than they would have taken to do correctly in the first place.

The counterintuitive approach works better: start manually.

Collect the first hundred records by hand. Categorize them yourself. Write the summaries yourself. Assign the tags yourself. Do the work before you automate the work.

This matters for two reasons. First, you can't design a good prompt pipeline for a process you don't fully understand. The worker prompts need to reflect how the categorization actually works in practice — the edge cases, the records that don't fit neatly, the decisions that require genuine judgment. You only learn those things by doing the process manually enough times to see the patterns.

Second, the manual phase forces you to design your taxonomy before you scale it. A database where the first 100 records were categorized carefully and consistently is a foundation you can build on. A database where the first 10,000 records were auto-categorized against a half-thought-through taxonomy is a problem you'll spend months cleaning up.

Understand the process first. Then let AI handle the scale.

The Interface Can Be Primitive

This is one of the most liberating truths about building intelligence products: users do not care about your interface. They care about finding the answer they came for.

A plain, functional interface built around genuinely exceptional data will outperform a beautifully designed interface built around weak data every single time. People are not paying for aesthetics. They're paying for better decisions faster.

See also The Difference Between a Prompt and a System Prompt (And Why It Matters for Your Business)

This means the barrier to shipping is lower than it appears. An Airtable base with good filtering. A simple directory built on a no-code tool. A searchable spreadsheet made public. A basic web interface with a clean search bar. None of these are impressive technically. All of them are sufficient to deliver value if the intelligence underneath them is real.

The design can improve over time. The data is what determines whether there's

anything worth improving the design for.

Your Existing Work May Already Be the Foundation

This is worth examining honestly if you've been creating content, building tools, or developing a prompt library in any capacity.

A prompt library is already a database. With the right categorization system — by use case, industry, output type, complexity, framework — it becomes a prompt intelligence platform that serious users would pay to access rather than a free resource they visit once.

A blog archive is already a research repository. Organized by topic, linked intelligently, supplemented with the source material behind each piece — it becomes an industry knowledge hub rather than a reverse-chronological list of posts.

An AI tools directory is already market intelligence. With pricing data, integration information, competitor mapping, and regular updates — it becomes a platform that product managers, investors, and operators reference professionally.

The pieces may already exist. The opportunity is often in connecting them deliberately and building the organizational layer that turns a collection into a system.

The Model in Plain Terms

You don't need to build software first. You need to build intelligence first.

Identify one valuable question a specific audience genuinely needs answered. Design the organizational architecture that structures the answer at scale — the taxonomy, the categories, the scoring criteria, the cross-references. Build the worker prompt pipeline that lets you collect and process records efficiently. Start manually, understand the process, then automate what you understand. Add records consistently, every week, without stopping. Let the database compound.

At some point — later than feels comfortable, earlier than you expect — the intelligence asset becomes valuable enough that people will pay to access it. At that point, the interface you build around it is almost secondary. What you've actually built is something scarce in a world drowning in generated content: organized, maintained, expert-curated intelligence on a topic that matters to people who make decisions with it.

That's the Micro SaaS. Not the software. The intelligence.

The software is just how people open the door.

Browse the prompt library at theronclaude.com →

Share this:

- Share on Facebook (Opens in new window) Facebook
- Share on X (Opens in new window) X